

WHAT IS CLAIMED IS:

1. A method of interpreting a linguistic representation of a textual input, comprising:
  - receiving a linguistic discourse representation structure (DRS) as the linguistic representation of the textual input;
  - receiving an entity-and-relation model of a non-linguistic domain; and
  - generating a semantic discourse representation structure (SemDRS) in terms of the entity-and-relation model and based on evidence derived from the linguistic DRS.
2. The method of claim 1 wherein generating a SemDRS comprises:
  - receiving a set of semantic mapping rules, each rule having a first side that matches DRS segments of a specified form and a second side that specifies at least a partial SemDRS; and
  - applying the set of semantic mapping rules to the linguistic DRS.
3. The method of claim 2 wherein generating a SemDRS comprises:
  - identifying applicable semantic mapping rules as rules having first sides that match any segments of the linguistic DRS; and
  - generating patterns associated with the linguistic DRS, the patterns including a

plurality of partial SemDRSs corresponding to the second sides of the applicable semantic mapping rules.

4. The method of claim 3 wherein generating a SemDRS comprises:

combining some of the patterns, consistently with the entity-and-relation model, into a single combined solution pattern.

5. The method of claim 4 wherein the linguistic DRS has elements arranged in a box structure and wherein generating a SemDRS further comprises:

arranging the solution pattern in a box structure based on the box structure of the linguistic DRS.

6. The method of claim 5 wherein arranging the solution pattern in a box structure comprises:

restoring the box structure of the linguistic UDRS onto the solution pattern to obtain the SemDRS.

7. The method of claim 5 and further comprising:  
prior to combining some of the patterns,

generating at least one initial search state including a set of the patterns, wherein the set of patterns is formed such that the patterns in the set account for every element in the linguistic DRS box

structure and such that none of the elements in the linguistic DRS box structure is accounted for by more than one pattern in the set of patterns.

8. The method of claim 7 wherein combining some of the patterns comprises:

selecting an initial search state; and  
combining all patterns in the selected initial search state into one combined pattern consistently with the entity-and-relation model.

9. The method of claim 8 wherein generating a SemDRS comprises:

for each pattern, generating a pat2drselems mapping between elements in the linguistic DRS and the patterns based on the element of the linguistic DRS represented by the first side of the applicable semantic mapping rule that spawned the pattern, and wherein generating at least one initial search state includes accessing the pat2drselems mapping.

10. The method of claim 9 wherein generating at least one initial search state comprises:

identifying a set of preferences corresponding to the set of patterns, the preferences

being indicative of how the patterns are to be combined to obtain the solution pattern.

11. The method of claim 10 and further comprising:  
generating the set of preferences.
12. The method of claim 11 wherein generating patterns comprises:  
generating a set of semantic entities or semantic relations, each relation having an associated source entity and a destination entity;  
loosening the patterns to obtain loosened patterns by replacing at least some of the source and destination entities in the set of patterns with abstract entities and wherein generating the set of preferences includes generating at least one preference for each loosened pattern indicative of how the loosened pattern is to relate to other patterns.
13. The method of claim 7 wherein generating patterns comprises:  
identifying elements in the linguistic DRS that are not associated with an applicable semantic mapping rule; and  
generating a semantic entity pattern corresponding to each of the identified elements.

14. The method of claim 13 wherein generating the initial search state comprises:

identifying words in the textual input that have a semantic content represented by the box structure of the linguistic DRS but that are not associated with a pattern; and generating an ignorable pattern corresponding to each of the identified words.

15. The method of claim 14 and further comprising: receiving a list of tokens in the textual input; maintaining a mapping of the tokens to the elements of the linguistic DRS; and maintaining a mapping between elements of the linguistic DRS and the patterns generated.

16. The method of claim 14 wherein generating a SemDRS further comprises:

identifying tokens that have a semantic content represented by the box structure of the linguistic DRS; applying string-based rewrite rules to the identified tokens; and generating additional patterns based on application of the string-based rewrite rules.

17. The method of claim 16 and further comprising:

prior to applying string-based rewrite rules to the identified tokens; and  
generating additional patterns, determining whether sufficient solution patterns have been generated and only applying the string-based rewrite rules to the identified tokens and generating the additional patterns if insufficient solution patterns have been created.

18. The method of claim 17 wherein generating a SemDRS further comprises:

generating initial search states based on the additional patterns.

19. The method of claim 18 wherein generating a SemDRS further comprises:

ranking the initial search states created from patterns generated by applying the semantic mapping rules higher than initial search states created from the additional patterns.

20. The method of claim 1 wherein generating a SemDRS comprises:

generating a plurality of SemDRSs ranked in order of cost.

21. A method of generating a semantic interpretation of a textual input represented by a linguistic discourse representation structure (DRS), comprising:

applying semantic mapping rules to the

linguistic DRS, mapping portions of the linguistic DRS to semantic interpretation fragments;

applying string-based rewrite rules to tokens in the textual input, mapping tokens in the textual input to semantic interpretation fragments; and

generating a plurality of ranked semantic discourse representation structures (SemDRSs) based on the semantic interpretation fragments, ranking SemDRSs generated from interpretations spawned by applying the semantic mapping rules higher than SemDRSs based on semantic interpretation fragments spawned by applying the string-based rewrite rules.

22. The method of claim 21 wherein the semantic interpretation fragments correspond to portions of an entity-and-relation model of a non-linguistic domain.

23. The method of claim 22 wherein the entity-and-domain model comprises an application schema.

24. The method of claim 22 wherein generating ranked SemDRSs comprises:

setting up a set of initial search states, each initial search state including a plurality of patterns such that each element of the linguistic DRS is covered by one and only one pattern.

25. The method of claim 24 wherein generating ranked SemDRSs comprises:

combining the patterns in each initial search state until the patterns in at least one initial search state are combined into a single solution pattern.

26. The method of claim 25 wherein setting up a set of initial search states includes identifying a set of preferences corresponding to each initial search state, the preferences indicating how the patterns in the initial search state relate to one another.

27. The method of claim 26 wherein combining the patterns comprises combining the patterns according to the preferences.

28. The method of claim 27 wherein the linguistic DRS has a box structure and wherein generating ranked SemDRSs comprises:

applying the box structure of the linguistic DRS to the solution pattern.

29. A semantic analysis system, comprising:



a controller configured to receive a linguistic discourse representation structure (DRS) as a linguistic representation of a textual input, an entity-and-relation model of a non-linguistic domain, and a set of semantic mapping rules;

an interpretation fragment generator, coupled to the controller, configured to apply the semantic mapping rules to the linguistic DRS to generate semantic interpretation fragments; and

an interpretation assembly component, coupled to the controller, receiving the semantic interpretation fragments and generating at least one solution pattern from the semantic interpretation fragments, the controller generating a semantic discourse representation structure (SemDRS) from the solution patterns consistent with the entity-and-relation model.

30. The semantic analysis system of claim 29 wherein each of the semantic mapping rules has a first side that matches DRS segments of a specified form and a second side that specifies at least a partial SemDRS.

31. The semantic analysis system of claim 30 wherein the interpretation assembly component generates the solution pattern by combining some of the interpretation fragments, consistently with the

entity-and-relation model, into a single combined solution pattern.

32. The semantic analysis system of claim 31 wherein the linguistic DRS has elements arranged in a box structure and wherein the controller is configured to generate the SemDRS by arranging the solution pattern in a box structure based on the box structure of the linguistic DRS.

33. The semantic analysis system of claim 32 wherein the controller comprises:

- a search state generator configured to generate at least one initial search state including a set of interpretation fragments, wherein the set of interpretation fragments is formed such that the interpretation fragments in the set account for every element in the linguistic DRS box structure and such that none of the elements in the linguistic DRS box structure is accounted for by more than one interpretation fragment.

34. The semantic analysis system of claim 33 wherein the interpretation fragment generator comprises:

- a blurb processor configured to identify elements in the linguistic DRS that are to be possibly interpreted as strings and to generate a semantic entity interpretation

fragment corresponding to each of the identified elements.

35. The semantic analysis system of claim 34 wherein the controller receives a list of tokens in the textual input and is configured to maintain a mapping of the tokens to the elements of the linguistic DRS and a mapping between elements of the linguistic DRS and the interpretation fragments generated.

36. The semantic analysis system of claim 35 wherein the interpretation fragment generator comprises:

a string-based fragment generator configured to identify tokens that have a semantic content represented by the box structure of the linguistic DRS, apply string-based rewrite rules to the identified tokens, and generate additional interpretation fragments based on application of the string-based rewrite rules.

37. The semantic analysis system of claim 36 wherein the search state generator is configured to generate initial search states based on the additional interpretation fragments.

38. The semantic analysis system of claim 37 wherein the search state generator is configured to rank the initial search states created from interpretation fragments generated by applying the semantic mapping

rules higher than initial search states created from the additional interpretation fragments.

39. A control component in a semantic analysis system configured to maintain a plurality of data structures for semantically interpreting a linguistic discourse representation structure (DRS) that is a linguistic representation of a textual input having tokens, the data structures comprising:

- a pattern list including a list of semantic patterns generated by applying semantic mapping rules to structural elements of the linguistic DRS;
- a pattern-to-DRS element mapping that maps patterns in the pattern list to DRS elements that spawned the patterns;
- a DRS element-to-token mapping that maps between the structural elements of the linguistic DRS and the tokens in the textual input that spawned the structural elements; and
- a pattern-to-token mapping between the patterns in the pattern list and tokens corresponding to the DRS elements used in generating the patterns.

40. The control component of claim 39 and further comprising:

- a preference list including preferences for how the patterns in the pattern list relate to other patterns in the pattern list.

41. The control component of claim 39 wherein the pattern list further includes semantic rewrite patterns generated by applying string-based rewrite rules to the tokens in the textual input.

42. The control component of claim 41 wherein the semantic rewrite patterns are marked to distinguish them from other patterns in the pattern list.

43. The control component of claim 39 and further comprising:

an entity-and-relation model for a non-linguistic domain.

44. The control component of claim 43 and further comprising:

a set of the semantic mapping rules that map between elements of the linguistic DRS and portions of the entity-and-relation model.